

REMARKS

Claims 1-18 and 20-136 are pending in the application. All claims have been rejected. These rejections are respectfully traversed and reconsideration is requested. The claims have been amended to more clearly point out what the Applicants regard as patentable subject matter and are not related to reasons of patentability. All claims are now believed to be in condition for allowance.

Claim Rejections Under 35 U.S.C. § 102

Claims 27-38, 44-46, and 48-52 have been rejected under 35 U.S.C. § 102(e) as being anticipated by Adya (U.S. Patent No. 6,266,658). Claim 27 is an independent claim. This rejection is respectfully traversed and reconsideration is requested.

The claims are directed to a method for evaluating a plurality of candidate index sets for a workload of database statements in a database system. An index superset is formed from a union of a current index set and a proposed index set. The current index set is a set of indexes that have been created for the database. The proposed index set may be proposed by specialized tools, such as Oracle Expert, or a user, such as an experienced database administrator. One or more candidate index sets are derived from this superset and are included or incorporated into a pool of candidate index sets. Additional candidate index sets may be derived from the pool of previously derived candidate index sets which may in turn be included or incorporated into the pool. Statistics are generated based on the plurality of derived candidate index sets and presented to a user or an index tuning mechanism.

Adya discusses a technique for identifying a single set of database indexes given a workload for the database. (Col. 3, ll. 23-26). According to Adya, a workload containing the most common queries to the database is parsed to identify candidate indexes that may be useful to satisfy the queries. (Col. 6, ll. 61-64). A plan of execution is derived from the queries and the set of candidate indexes. This plan contains index usage information that identifies which indexes are used to answer each query. (Col. 6, l. 55 to col. 7, l. 11). The candidate indexes are further evaluated on the basis of the cost of the plan and the indexes with the most potential for improvement of the cost are selected. (Col. 7, ll. 12-20). The selected indexes are further

evaluated to ensure that the total workload cost never increases with low benefit. (Col. 7, ll. 22-27 and col. 9, ll. 2-21). The outcome of this evaluation is a single set of indexes that may be used by a user to create new indexes and modify the set of indexes for use on the database to enhance performance. (Col. 7, ll. 30-33).

The Applicants respectfully submit that Adya does not teach or suggest including or incorporating the derived candidate index sets into a pool of candidate index sets from which additional candidate index sets may be derived, as is recited, for example, in independent Claim 27. Instead, Adya discusses forming a candidate index set by parsing queries associated with a workload to identify only a single set of potential indexes. In contrast to Adya, the Applicants claim deriving a plurality of candidate index sets from a superset that is formed by the union of a current index set and a proposed index set, and incorporating the derived candidate index sets into a pool of candidate index sets from which additional candidate index sets may be derived. Nowhere does Adya teach or suggest placing the derived candidate index sets into a pool of candidate index sets.

Applicants therefore respectfully request reconsideration and withdrawal of the rejection of amended Claim 27 in favor of allowance. Dependent Claims 28-38, 44-46, and 48-52 depend on independent Claim 27 and are not anticipated by Adya for at least the same reasons as above.

Claim Rejections Under 35 U.S.C. § 103

Claims 1-6, 8-12, 16, 18, 20, 22-24, 26, 53-54, 56, 58-60, 62-64, 66-70, 74, 76-77, 79-81, 83-86, 88-90, 92-96, 100, 102-103, 105-107, 109-113, 115-117, 119-123, 127, 129-130, 132-134, and 136 have been rejected under 35 U.S.C. § 103(a) as being unpatentable over Chaudhuri I (U.S. Patent No. 5,960,423) in view of Chaudhuri II (U.S. Patent No. 6,223,171).

Claims 7, 65, 91, and 118 have been rejected under 35 U.S.C. § 103(a) as being unpatentable over Chaudhuri I and Chaudhuri II in further view of Jakobsson (U.S. Patent No. 5,924,088).

Claims 13-15, 71-73, 97-99, and 124-126 have been rejected under 35 U.S.C. § 103(a) as being unpatentable over Chaudhuri I and Chaudhuri II in further view of Eberhard (U.S. Patent No. 6,003,022).

Claims 39-41 have been rejected under 35 U.S.C. § 103(a) as being unpatentable over Adya in further view of Eberhard.

Claims 17, 75, 101, and 128 have been rejected under 35 U.S.C. § 103(a) as being unpatentable over Chaudhuri I and Chaudhuri II in further view of Smith (U.S. Patent No. 5,404,510).

Claim 43 has been rejected under 35 U.S.C. § 103(a) as being unpatentable over Adya in further view of Smith.

Claims 21, 78, 104, and 131 have been rejected under 35 U.S.C. § 103(a) as being unpatentable over Chaudhuri I and Chaudhuri II in further view of Celis (U.S. Patent No. 6,021,405).

Claim 47 has been rejected under 35 U.S.C. § 103(a) as being unpatentable over Adya in further view of Celis.

Claims 25, 82, 108, and 135 have been rejected under 35 U.S.C. § 103(a) as being unpatentable over Chaudhuri I and Chaudhuri II in further view of Adya.

Claims 55, 61, 87, and 114 have been rejected under 35 U.S.C. § 103(a) as being unpatentable over Chaudhuri I and Chaudhuri II in further view of Gurry (“Oracle Performance Tuning”).

Claim 57 has been rejected under 35 U.S.C. § 103(a) as being unpatentable over Adya in further view of Gurry.

Claim 42 has been rejected under 35 U.S.C. § 103(a) as being unpatentable over Adya in further view of Finkelstein (“Physical Database Design for Relational Database”).

Claims 1, 53, 54, and 110 are independent claims. All of the above rejections are respectfully traversed and reconsideration is requested.

Similar to Claim 27, discussed above, independent Claims 1, 53, 54, and 110 recite incorporating the derived candidate index sets into a pool of candidate index sets from which additional candidate index sets may be derived. Claims 53 and 54 further recite that the candidate index sets are repeatedly derived.

Chaudhuri I discusses a technique for determining a single set of candidate indexes for a workload of queries that are to be executed against a database. (Col. 2, ll. 13-14). According to

the technique, index configurations are selected based on a relative effectiveness among candidate index configurations for the database. The effectiveness of a particular index configuration is based on cost estimates to execute queries of the workload against the database using the particular index configuration. (Col. 6, ll. 39-45). A total cost of an index configuration is determined as a sum of cost estimates to execute each query of the workload against the database using the particular index configuration. The candidate index configuration having the least total cost is considered to be the most effective in minimizing the cost of executing the workload against the database. (Col. 6, ll. 46-54).

Chaudhuri I also discusses using what-if indexes to simulate the presence of indexes absent from the database. (Col. 7, ll. 11-15). These what-if indexes are used to evaluate candidate index configurations. (Col. 8, ll. 28-41).

Chaudhuri II describes a what-if analysis technique that may be used to perform quantitative analysis of a database system. (Col. 1, ll. 64-65). A hypothetical configuration is defined for the database system. The hypothetical configuration contains a hypothetical set of indexes and is of a hypothetical size. (Col. 2, ll. 21-24, col. 13, ll. 24-26, and Fig. 5). The hypothetical configuration is simulated and the costs of workload queries are estimated with respect to the hypothetical configuration. (Col. 2, ll. 26-44, col. 14, ll. 23-24, and Fig. 5). The evaluation continues until all queries for the workload have been evaluated. (Col. 15, ll. 18-23, and Fig. 5). After all queries for the workload have been evaluated, an estimated cost for the workload is determined from the estimated costs of all the queries. (Col. 15, ll. 53-57, and Fig. 5).

Although Chaudhuri I and Chaudhuri II may evaluate candidate index configurations, the Applicants respectfully submit that Chaudhuri I and Chaudhuri II, taken either separately or in combination, do not teach or suggest incorporating the derived candidate index sets into a pool of candidate index sets from which additional candidate index sets may be derived. Furthermore, the cited references do not suggest repeatedly deriving the candidate index sets. Instead, they merely discuss selecting a single set of candidate indexes for a workload of queries.

Applicants, therefore, respectfully request reconsideration and withdrawal of the rejections of Claims 1, 53, 54, and 110 as amended, in favor of allowance. Allowance of dependent Claims 2-6, 8-12, 16, 18, 20, 22-24, 26, 56, 58-60, 62-64, 66-70, 74, 76-77, 79-81, 83-

86, 88-90, 92-96, 100, 102-103, 105-107, 109, 111-113, 115-117, 119-123, 127, 129-130, 132-134, and 136 follows from the allowance of Claims 1, 53, 54, and 110.

Jakobsson discusses a method for generating an index access path from a list of indexes and a “where” clause. Jakobsson does not supplement the failure of Chaudhuri I and Chaudhuri II to teach or suggest incorporating the derived candidate index sets into a pool of candidate index sets from which additional candidate index sets may be derived, and, therefore, Claims 7, 65, 91, and 118 are not rendered obvious by the combination of Chaudhuri I, II and Jakobsson.

Eberhard discusses a method for estimating the execution costs of an SQL statement, transaction, or application when accessing a database. Eberhard does not supplement the failure of Chaudhuri I and Chaudhuri II to teach or suggest incorporating the derived candidate index sets into a pool of candidate index sets from which additional candidate index sets may be derived, and therefore Claims 13-15, 71-73, 97-99, and 124-126 are not rendered obvious by the combination of Chaudhuri I, II and Eberhard.

Eberhard does not supplement the failure of Adya to teach or suggest incorporating the derived candidate index sets into a pool of candidate index sets from which additional candidate index sets may be derived, and therefore Claims 39-41 are not rendered obvious by the combination of Adya and Eberhard.

Smith discusses a method for establishing indexes which provide for efficient operation of a database system based on an analysis of the workload with weight given to the importance of specific requests. Smith does not supplement the failure of Chaudhuri I and Chaudhuri II to teach or suggest incorporating the derived candidate index sets into a pool of candidate index sets from which additional candidate index sets may be derived, and therefore Claims 17, 75, 101, and 128 are not rendered obvious by the combination of Chaudhuri I, II and Smith.

Smith does not supplement the failure of Adya to teach or suggest incorporating the derived candidate index sets into a pool of candidate index sets from which additional candidate

index sets may be derived, and therefore Claim 43 is not rendered obvious by the combination of Adya and Smith.

Celis discusses a method for optimizing SQL database queries that includes a “tracking mechanism” which eliminates the generation of redundant expressions. Celis does not supplement the failure of Chaudhuri I and Chaudhuri II to teach or suggest incorporating the derived candidate index sets into a pool of candidate index sets from which additional candidate index sets may be derived, and therefore Claims 21, 78, 104, and 131 are not rendered obvious by the combination of Chaudhuri I, II and Celis.

Celis does not supplement the failure of Adya to teach or suggest incorporating the derived candidate index sets into a pool of candidate index sets from which additional candidate index sets may be derived, and therefore Claim 47 is not rendered obvious by the combination of Adya and Celis.

Adya does not supplement the failure of Chaudhuri I and Chaudhuri II to teach or suggest incorporating the derived candidate index sets into a pool of candidate index sets from which additional candidate index sets may be derived, and therefore Claims 25, 82, 108, and 135 are not rendered obvious by the combination of Chaudhuri I, II and Adya.

Gurry discusses the use of statistics to determine whether an index in a database is a candidate for dropping and recreating. Gurry does not supplement the failure of Chaudhuri I and Chaudhuri II to teach or suggest incorporating the derived candidate index sets into a pool of candidate index sets from which additional candidate index sets may be derived, and therefore Claims 55, 61, 87, and 114 are not rendered obvious by the combination of Chaudhuri I, II and Gurry.

Gurry does not supplement the failure of Adya to teach or suggest incorporating the derived candidate index sets into a pool of candidate index sets from which additional candidate

index sets may be derived, and therefore Claim 57 is not rendered obvious by the combination of Adya and Gurry.

Finkelstein discusses a physical design tool for relational databases that suggests a physical configuration given a workload consisting of a set of SQL statements and their execution frequencies. Finkelstein does not supplement the failure of Adya to teach or suggest incorporating the derived candidate index sets into a pool of candidate index sets from which additional candidate index sets may be derived, and therefore Claim 42 is not rendered obvious by the combination of Adya and Finkelstein.

All claims are now believed to be in condition for allowance.

CONCLUSION

In view of the above amendments and remarks, it is believed that all claims are in condition for allowance, and it is respectfully requested that the application be passed to issue. If the Examiner feels that a telephone conference would expedite prosecution of this case, the Examiner is invited to call the undersigned attorney.

Respectfully submitted,

HAMILTON, BROOK, SMITH & REYNOLDS, P.C.

By Lyudmila Lubashev
Lyudmila Lubashev
Registration No. 55,408
Telephone: (978) 341-0036
Facsimile: (978) 341-0136

Concord, MA 01742-9133
Dated: 7/16/2005